LEVEL II

AD A071604

DDC FILE COPY

UNIVERSITY OF MARYLAND

COMPUTER SCIENCE CENTER

COLLEGE PARK, MARYLAND

20742

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

LEVEL II

(9) Technical rept.,

(14) TR-707
DAAG-53-76C-0138

(11) October, 1978

(6) RELAXATION APPLIED TO MATCHING QUANTITATIVE RELATIONAL STRUCTURES.

(12) 26p.

(10) Les Kitchen
Computer Science Center
University of Maryland
College Park, MD 20742

ABSTRACT

A relaxation process is applied to the problem of matching relational structures involving numeric quantities. The method is empirically shown to be very sensitive to the exact form of the updating rule used. However, with the proper updating rule, the method works well, and is remarkably tolerant of measurement errors.

1473

(15) The support of the Defense Advanced Research Projects Agency and the U.S. Army Night Vision Laboratory under Contract DAAG-53-76C-0138, (16) DARPA order-3206 is gratefully acknowledged, as is the help of Kathryn Riley in preparing this paper. The author would also like to thank Dr. Azriel Rosenfeld for many stimulating discussions.

403 018

JB

1. Introduction

This paper addresses the following problem: Suppose we have a description of a certain object in terms of its parts, their properties, and the relations between them. Now suppose we are given a description of some part of the real world in terms of these same properties and relations. How can we determine whether an instance of this object actually occurs in the given part of the world? A good example is the interpretation of visual images. We may describe a house, for example, in terms of its parts (roof, walls, windows, doors), their properties (such as shape, orientation, color and texture), and the spatial relations between them (roof above walls, windows surrounded by walls, etc.). Given a visual scene, we may segment its image into regions, and describe the scene in terms of these regions, their properties and their relations. We may now ask the question, Is there a house in this scene? That is, are there regions in the scene which look like roofs, walls, windows etc., and are those regions in the proper spatial relationships to form a house?

A previous paper (Kitchen 1978) has dealt with descriptions using properties and relations that must either be wholly true or wholly false. However, these are not fully adequate for describing real-world objects, so we are led to consider properties and relations which can take on numeric values, and along with this, we must be prepared to deal with errors in

such values, permitting "approximate" instances of a description of an object. The next section provides a formalism for such "quantitative" properties and relations, and indicates how relaxation methods can be applied to matching descriptions in these terms. The third section contains an empirical study of the effectiveness of these methods, while the fourth discusses the results of this study and ways in which relaxation methods can be further extended.

2. Definitions and Description of Method

The following provides a formal notation for descriptions of an object (or a scene) in terms of its parts, their properties, and the relations between them.

Definition

A quantitative relational structure is a triple $\langle X,\Phi,\varphi\rangle$ whose components are as follows:

X is a finite set, called the carrier of the relational structure. The elements of X are called nodes, and correspond to the atomic parts of an object which are used in its description.

Φ is a finite set of quantities. Each quantity in Φ has associated with it a pair of positive integers, called respectively its order and dimensionality, indicated by superscripts in parentheses. Thus $Q^{(m,n)} \in \Phi$ indicates that Q is a quantity in Φ with order m and dimensionality n. For example, a quantity corresponding to the ordinary Cartesian coordinates of a point would have order 1 and dimensionality 2, since it assigns to single points pairs of numbers. On the other hand, a quantity corresponding to the ordinary Euclidean distance between points would be of order 2 and dimensionality 1, since it assigns a single number to each pair of points.

φ is a mapping that associates with each $Q^{(m,n)} \in \Phi$ a partial

function from X^m to $\underline{R}^n$, where $\underline{R}$ is the set of real numbers.

Thus φ determines exactly which m-tuples of nodes have values associated with them by quantity Q, and what these values are.

Let $Q^{(m,n)} \in \Phi$, and $x_1, x_2, \ldots, x_m \in X$. By way of notation we write $Q^{(m,n)}(x_1, x_2, \ldots, x_m)$ to indicate the value of $\varphi(Q^{(m,n)})$ at $(x_1, x_2, \ldots, x_m)$, provided the function is defined at this point. Following the usual convention, when no confusion can arise, a quantitative relational structure will be denoted merely by the name of its carrier.

Below, the symbol M will normally be used to denote a quantitative relational structure thought of as a model, that is, a description of a certain object in terms of its parts and their properties. Similarly, the symbol W will normally be used to denote a quantitative relational structure which describes some part of the world.

In order to find an instance of an object in some part of the world, it is necessary to match up the parts of the object with the corresponding pieces in the world. However, because of errors and uncertainties in measurements, we must be prepared to assign different degrees of confidence to different pairings of model parts with world parts. It therefore becomes necessary to introduce the notions of fuzzy truth and fuzzy logic. In ordinary logic, a given proposition can be assigned one of two truth-values: it can be either "true" or "false".

The usual logical connectives such as "and" and "or" can be thought of as truth-functions which take as arguments either "true" or "false", and return either "true" or "false" as results, according to the truth-table for each. In the most common treatment of fuzzy logic, a given proposition can have assigned to it a real number in the range zero to one, reflecting the degree of truth in that proposition. Similarly, the truth-functions are extended to take real-valued arguments and produce real-valued results. Suitable extensions of the "and" function are xy or $\min\{x,y\}$ for arguments x and y. For the "or" function we can use $\max\{x,y\}$ or $1-(1-x)(1-y)$. For a more detailed treatment the reader is referred to Zadeh (1965).

Definition

Let M and W be quantitative relational structures. An assignment of M to W (the order being material) is a function which assigns to every pair $(x,y) \in M \times W$ a real number in the range zero to one. Thus an assignment indicates the degree of confidence we have in each pairing of model node with world node.

Definition

Let M and W be quantitative relational structures. Let $Q^{(m,n)}$ be a quantity. Then the goodness of fit function for Q, written γ_Q, is a function $\gamma_Q : \underline{R}^n \times \underline{R}^n \to [0,1]$. If $u \in \underline{R}^n$ is a value of Q in M, and $v \in \underline{R}^n$ a value of Q in W, then $\gamma_Q(u,v)$

measures how closely u and v agree, with 1.0 indicating perfect agreement, and 0.0 indicating no agreement whatsoever. Clearly, the choice of γ_Q is determined by our tolerance for measurement errors. A suitable form for γ_Q is:

$$\gamma_Q(u,v) = \frac{1}{1+\sum_{i=1}^{n}\left(\frac{(u_i-v_i)}{s_i}\right)^2}$$

where u_i and v_i represent the ith components of u and v respectively, and $s_1, s_2, \ldots, s_n$ are scaling parameters associated with Q. This form of the goodness of fit function will be the only one used in this paper. It has a peak of 1.0 when u and v are exactly equal, and falls off towards zero as u and v move apart.

<u>Definition</u>

Let M and W be quantitative relational structures with quantities Φ, and let A be an assignment of M to W. Let $x \in M$ and $y \in W$. Then the <u>local compatibility</u> of x with y, with respect to A, C(x,y,A) is given by

$$\bigcap_{\substack{x_1,x_2,\ldots,x_m \in M \\ Q^{(m,n)} \in \Phi \\ u \in \underline{R}^n \\ \text{such that} \\ x=x_i \text{ for some } i \\ Q(x_1,\ldots,x_m)=u}} \left[\bigcup_{\substack{y_1,y_2,\ldots,y_m \in W \\ v \in \underline{R}^n \\ \text{such that} \\ y=y_j \text{ for some } j \\ Q(y_1,\ldots,y_m)=v}} \left[\left(\bigcap_{1 \le k \le m} A(x_k,y_k) \right) \cap \gamma_Q(u,v) \right] \right]$$

where $\cap$ and $\cup$ stand for the operations of fuzzy logical "and" and fuzzy logical "or" respectively. While this formula may seem rather daunting at first sight, it falls naturally into three parts, each of which has a rather simple justification.

The innermost part

$$(\bigcap_{1\le k\le m} A(x_k,y_k))\cap\gamma_Q(u,v)$$

captures the following notion: Suppose we have two instances of a quantity, $Q(x_1,\ldots,x_m)=u$ and $Q(y_1,\ldots,y_m)=v$. Then these two instances can only be considered "compatible" if their numeric values match (hence the terms $\gamma_Q(u,v)$) <u>and</u> the assignment permits all the pairings $(x_1,y_1),(x_2,y_2),\ldots,(x_m,y_m)$ (hence the term $\bigcap_{1\le k\le m} A(x_k,y_k)$).

Moving outwards, the fuzzy logical "or" $\cup$, embodies the following: Suppose we have an instance of a quantity $Q(x_1,\ldots,x_m)$ which mentions x. Then in order for x to be "compatible" with y, there must be in the world a corresponding instance $Q(y_1,\ldots,y_m)$ which mentions y, and which is "compatible" in the sense of the previous paragraph.

Lastly, the outermost fuzzy logical "and" $\cap$, ensures that this "compatibility" is satisfied for every quantity instance in the model which mentions x.

Note that there is a different local compatibility function for each choice of fuzzy logical truth functions.

3. Relaxation

With the above ground-work, we can define a relaxation scheme which generates a sequence $A_0, A_1, A_2, \ldots$ of assignments between two quantitative relational structures M and W, using the following recurrence:

$A_0(x,y) = 1.0$ for all $x \in M$, $y \in W$

$A_i(x,y) = C(x,y,A_{i-1})$ for all $x \in M$, $y \in W$, $i>0$.

That is, we set up an initial assignment A_0, which permits all pairings of model nodes with world nodes. Then we use the local compatibility of each pair to generate a new assignment. This process can be repeated for a fixed number of steps, or until the sequence of assignments converges in some suitable sense. In fact, if we use maximum as a fuzzy logical "or", and minimum as a fuzzy logical "and" in the local compatibility function, and provided we assume that goodness of fit functions are always greater than zero, then we can prove these two results as generalizations of those in Kitchen (1978):

(1) There exists j such that $A_j = A_k$ for all $k>j$.

(2) If we treat the M and W as discrete relational structures (that is, we ignore the values of quantities and regard them as discrete predicates) then the relation $R = \{(x,y) \mid A_j(x,y) > 0\}$ can be considered as a discrete assignment from M to W which contains all monomorphisms (structure preserving mappings) from M to W.

These results follow easily from those in Kitchen (1978) when it is realized that local compatibility is the fuzzy analog of the notion weak local consistency defined therein. Under conditions other than those stated above, no such results can be simply proved, and the experimental results described below suggest that no such results in fact hold.

4. Implementation and Empirical Results

A computer program was written in the programming language Pascal to implement the above relaxation scheme. However, two small changes were made. Firstly, the goodness of fit functions were modified to return zero if any term in the summation were greater than a certain bound. This was necessary to prevent arithmetic overflow when comparing widely disparate quantity values. In the experiments described here, this bound was set at 100. The other modification permitted the program to omit from further processing any pairing whose assignment value had fallen below a certain threshold. This saved computer time by avoiding re-computation for pairings with so little support that they could be disregarded. The threshold used here was 10^{-6}. In several trials, the use of such a threshold had a negligible effect on the course of the relaxation, but often reduced computation time by a factor of two or three.

Several experiments were run to investigate the behavior of the relaxation under various conditions. This section describes what was common to all experiments.

All the experiments used as a world a quantitative relational structure describing 44 cities of the eastern United States in terms of their populations and the road mileages between them. Thus "population" was a quantity of order 1

and dimensionality 1. which was defined at every city, taking values between 0 and 7,605,000. The scaling parameter s_1 in its goodness of fit function was 10,000. "Distance" was a quantity of order 2 and dimensionality 1, defined for every pair of cities with a direct road route between them. It took values between 37 and 387, and had a scaling parameter of 10. The data were drawn from the Rand-McNally 1976 Road Atlas.

Models of various sizes were generated by randomly selecting compact, connected groups of cities, and then including all populations of these cities and all direct distances between them. A perturbation of size d (say) could be introduced into a model by the following method. For every quantity value in the model, generate a random number uniformly distributed in the range [-d,+d], multiply it by the scaling parameter appropriate to that quantity, and add it to the quantity's value. Thus, a perturbation of size 10 would lead to an average deviation of 50 miles in distance measurements, which is roughly a 25% relative error. This feature permits investigation of how well the relaxation copes with errors in measurement.

The first experiment tried various forms of the local compatibility function under increasing perturbations of the models used. The different forms were as follows:

Form 1: maximum for fuzzy logical "or"
minimum for fuzzy logical "and"

Form 2: maximum for fuzzy logical "or"
product for fuzzy logical "and"

Form 3: complemented-product for fuzzy logical "or"
product for fuzzy logical "and"

The complemented-product has the formula

$1-(1-a)(1-b)$ for real arguments a and b.

Three other forms (4,5, and 6) of the local compatibility function were used. These corresponded to forms 1, 2 and 3, except that the outermost application of the fuzzy logical "and" in the local compatibility function was replaced by an arithmetic mean. Roughly speaking, these forms of the local compatibility function would be prepared to accept evidence that looks good on average, rather than insist that all evidence be found.

Two measurements were used to gauge the effectiveness of the relaxation. The first is called the "truth-value difference". In any assignment, there will be some pairings which are correct, and many pairings which are incorrect. Ideally, all correct pairings should receive the value one, all incorrect pairings zero. This ideal situation will not always occur, but a good measure of how well an assignment separates correct pairings from incorrect pairings can be obtained by taking the average truth-value assigned to correct pairings, and

subtracting from it the average truth-value of incorrect pairings. However, since we are interested in the separation of good pairings from bad, and not absolute truth values, the values used for computing the averages were not the raw truth-values of the assignment; instead they were linearly scaled by division by the largest truth-value in the assignment. This also permitted a more equitable comparison across the different forms of local compatibility, since all the truth-values were standardized. Note that this normalization occurs only when actually computing the truth-value difference; the assignment used for the relaxation was in no way altered. In the best case the truth-value difference will have value 1.0, at very worst it could be -1.0, however in the normal worst case it would take a value very close to zero, indicating that the assignment makes no significant distinction between correct and incorrect pairings.

The other measurement used was called "good separation". An assignment is said to have good separation when the correct pairings have higher truth value than any incorrect pairing. Since good separation is either present or absent, it cannot distinguish between those assignments which only just fail to have good separation and those which fail miserably. However it captures the important desideratum that the correct pairing of model with world can be found merely by selecting the best pairs in the assignment.

For each form of the local compatibility function, four randomly generated models of size 5 were used, and these models were perturbed by amounts 0 through 5, and 10. The relaxation process was run for five iterations, and measurements made on the final assignment. Figure 1 shows the average truth-value difference over the four models, while Figure 2 shows the number of models (out of the four) which exhibited good separation. The same series of random models and perturbations was used for each form of the local compatibility function, so that the results are comparable across the table (in each row, all six forms were applied to the very same matching problems). Figures 3, 4 and 5 apply only to form 1 of the local compatibility function, the best behaved of the six tried.

Figure 3 is derived from the same experiment as Figures 1 and 2. It shows the average truth-value difference at the end of each iteration. (The rightmost column of Figure 3 is identical with the leftmost column of Figure 1.) Each row of Figure 3 gives some idea of the convergence behavior of the relaxation process.

Figures 4 and 5 demonstrate the effect of changing the model size. For each of three sizes 5, 7, and 10, four models were generated of that size. Each model was perturbed by various amounts. Five iterations of the relaxation process

were applied, and measurements were taken on the final assignment. Figure 4 shows the truth-value differences (averaged over the four models). Figure 5 shows the number of models (out of four) which exhibited good separation. (The leftmost column of Figure 4 is identical with that of Figure 1; similarly for Figures 5 and 2.)

5. Discussion of Results and Extensions

One of the most striking things about the results shown in Figures 1 and 2 is that the ability of the relaxation process to tolerate noise depends very much on the particular form used for the local compatibility function. Those forms which use multiplication as a fuzzy logical "and" seem to fare the worst. While remarkable at first sight, this behavior has a rather simple explanation, which is best illustrated by an example: Suppose we combine three truth-values: 0.1, 0.1 and α, where $\alpha<0.1$. The minimum of these three numbers will reflect directly any changes in α; while with the product, any changes in α will be attenuated by a factor of $0.1 \times 0.1 = 0.01$. Thus, the use of multiplication will not produce such a good separation of truth-values in an assignment, even when a linear rescaling is done. In practice it is observed that while the truth values of incorrect pairings are reduced quickly to zero, the truth values of pairings without perfect support (but which are nonetheless almost correct) are also reduced very quickly to zero. However, when using minimum as a fuzzy logical "and", a pairing without perfect support is reduced to a non-zero level which reflects the strength of its weakest support, and stabilizes at that level. Thus multiplication works badly because it too readily throws out pairings which should have been retained.

On the other hand, the use of an average at the outermost level of the local compatibility function does not work as well as the use of a minimum there, because the relaxation then holds onto incorrect pairings which should have been discarded. It does this because some incorrect pairings may have generally good support, and only fail in one or two respects. To an average such pairs would look good, while a minimum would reject them. Note however that these remarks apply only to the experiments described above, where "noise" was introduced into a model by perturbing the values of quantities without otherwise altering the structure. It is also possible to introduce "noise" by deleting nodes and quantity instances from the world. In visual terms, this corresponds to identifying an object, parts of which may be obscured or otherwise invisible. For this type of problem, the averaging of evidence may very well be the better method, as is suggested by the work of Ranade and Rosenfeld (1978).

Barrow and Tenenbaum (1976) in their paper about MSYS remark that the particular choice of fuzzy logic functions seems not to have much effect on the relaxation. The different behavior observed here can probably be attributed to the fact that these experiments deal with larger matching problems, and with apparently larger perturbations than theirs.

In Figure 3 we see the improvement in the separation in the assignment. Not only does perturbation of the model degrade

the result at a given iteration, but it also increases the number of iterations required to reach convergence.

Figures 4 and 5 show that the size of the model has little effect on how good an answer the relaxation produces. It should be remarked in passing that the computation time varies in a roughly linear fashion with the size of the model. This is because the computation associated with each model node on a given iteration is fixed, depending only on the quantity instances which explicitly mention that node; while the number of iterations required for convergence seems unaffected by model size.

Relational structures can be extended in many ways to make them more useful tools. Some of these have already been mentioned (Kitchen 1978). However, one rather simple extension to quantitative relational structures would seem to be quite powerful. First of all, instead of having a general goodness of fit function for a quantity, we should permit a different goodness of fit function for each instance of a property in a model. By this means we could enforce a tighter match to more critical parts of a model, while easing the constraints on less critical parts. We could then make queries like "Find a city with exactly 68,000 people about 100 miles from another city with a population of roughly 100,000." Furthermore, there is no need for goodness of fit functions to be symmetric with

respect to deviations above and below the ideal. Using asymmetric functions we could pose problems such as "Find a town with population about 10,000 which is less than 100 miles from a city of at least 500,000 people."

Even without such embellishments, relaxation seems to be a very useful method for matching quantitative relational structures under distortion. With the use of a well-behaved local compatibility function, the method is remarkably tolerant of measurement errors.

References

1. Barrow, H.G. and Tenenbaum, J.M. (1976) "MSYS: A system for reasoning about scenes". Artificial Intelligence Center Technical Note 121, Stanford Research Institute, Menlo Park, CA.

2. Kitchen, L.J. (1978) "Discrete relaxation for matching relational structures". Computer Science Technical Report TR-665, University of Maryland, College Park, MD.

3. Ranade, S. and Rosenfeld, A. (1978) "Point pattern matching by relaxation". Computer Science Technical Report TR-702, University of Maryland, College Park, MD.

4. Zadeh, L.A. (1965) "Fuzzy sets". Information and Control, vol. 8, pp. 338-353.

	Form of local compatibility					
Perturbation	1	2	3	4	5	6
0	1.00	1.00	1.00	0.94	0.98	0.98
1	1.00	0.00	0.00	0.86	0.51	0.63
2	1.00	0.00	0.00	0.70	0.31	0.40
3	0.99	0.00	0.00	0.60	0.13	0.17
4	0.99	0.00	0.00	0.54	0.14	0.16
5	0.99	0.00	0.00	0.40	0.11	0.15
10	0.90	0.00	0.00	0.23	0.02	0.02

Figure 1. Truth-value difference for different forms of local compatibility

	Form of local compatibility					
Perturbation	1	2	3	4	5	6
0	4	4	4	4	4	4
1	4	0	0	4	1	1
2	4	0	0	1	0	1
3	4	0	0	0	0	0
4	4	0	0	0	0	0
5	4	0	0	0	0	0
10	0	0	0	0	0	0

Figure 2. Good separation for different forms of local compatibility

Perturbation	Iterations 1	2	3	4	5
0	0.98	1.00	1.00	1.00	1.00
1	0.79	0.95	1.00	1.00	1.00
2	0.57	0.78	0.94	1.00	1.00
3	0.33	0.72	0.96	0.99	0.99
4	0.32	0.73	0.92	0.99	0.99
5	0.23	0.83	0.98	0.99	0.99
10	0.10	0.53	0.86	0.86	0.90

Figure 3. Truth-value difference for form 1, by iteration

Perturbation	Model size 5	7	10
0	1.00	1.00	1.00
1	1.00	1.00	1.00
2	1.00	1.00	1.00
3	0.99	0.99	1.00
4	0.99	0.99	1.00
5	0.99	1.00	1.00
10	0.90	0.93	1.00

Figure 4. Truth-value difference for form 1 for various model sizes

	Model size		
Perturbation	5	7	10
0	4	4	4
1	4	4	4
2	4	4	4
3	4	4	4
4	4	4	3
5	4	3	2
10	0	3	2

Figure 5. Good separation for form 1 for various model sizes

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) RELAXATION APPLIED TO MATCHING QUANTITATIVE RELATIONAL STRUCTURES		5. TYPE OF REPORT & PERIOD COVERED Technical
		6. PERFORMING ORG. REPORT NUMBER TR-707
7. AUTHOR(s) Les Kitchen		8. CONTRACT OR GRANT NUMBER(s) DAAG-53-76C-0138
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Center, University of Maryland, College Park, MD 20742		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Night Vision Laboratory, Ft. Belvoir, VA 20060		12. REPORT DATE October, 1978
		13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Relational structures
Matching
Relaxation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A relaxation process is applied to the problem of matching relational structures involving numeric quantities. The method is empirically shown to be very sensitive to the exact form of the updating rule used. However, with the proper updating rule, the method works well, and is remarkably tolerant of measurement errors.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)